

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

UTILITY PATENT APPLICATION FOR:
LOAD MANAGEMENT IN A POWER SYSTEM

Inventors:

Keith Istvan Farkas
345 Laurel Street
San Carlos, CA 94070

Cullen Edwin Bash
662 Mangels Avenue
San Francisco, CA 94127

Parthasarathy Ranganathan
3375 Alma Street, #286
Palo Alto, CA 94306

HP Docket No. 200313156-1

LOAD MANAGEMENT IN A POWER SYSTEM

TECHNICAL FIELD

5 This invention relates generally to power systems. More particularly, the invention relates to managing the load on components in a power system.

BACKGROUND

10 Power systems typically include redundant components to prevent power outages. Figure 8 illustrates an example of using redundant components in a power system. An uninterruptible power supply (UPS) 810 and a UPS 820 supply power to a power distribution unit (PDU) 830, which may be connected to power supplies and electrical devices, i.e., the load, not shown. The UPS 810 and UPS 820 are redundant components and may equally
15 share the load demand of the PDU 830. That is the PDU 830 may draw a substantially equal amount of current from the UPS 810 and the UPS 820 to meet the load demand on the PDU 830. If the UPS 810 fails, then the UPS 820 is available to meet the load demand of the PDU 830 on its own. However, the UPS 820 must be sufficiently provisioned, i.e., must have sufficient capacity, such that the UPS 820 can meet the load demand of the PDU 830 on its
20 own. Furthermore, the UPS 820 may be sharing the load demand of another PDU 840 with another UPS (not shown), and if that UPS fails then the UPS 820 may have to support two loads on its own. Thus, the UPS 820 must be substantially over provisioned to prevent power outages.

 Over provisioning components in a power system and using redundant components
25 are fail-safe techniques, but these techniques dramatically increase operating and

implementation costs. Typically, as the capacity of a UPS or other power system component increases, costs also increase. Furthermore, even if redundancy and over provisioning are used, the power system may still be susceptible to the “snow ball effect”. For example, if the UPS 810 fails then the UPS 820 must meet the load demand of the PDUs 830 and 840. At
5 peek demand intervals, the combined demand of the PDUs 830 and 840 may exceed the capacity of the UPS 820. For example, the UPS 820 may have been designed to have a capacity to meet a lower peek demand, but the peek demand may have increased over time. If the UPS 820 fails, the resulting increased loading may cause other UPSs in the power system to fail, realizing the snow ball effect.

10

SUMMARY OF THE EMBODIMENTS

According to an embodiment, a method of managing load in a power system comprises determining whether a load demand on at least one power system component of a plurality of power system components needs to be varied. The method further comprises
15 determining a new load demand to be placed on the power system component in response to determining the load demand on the power system component needs to be varied, wherein the new load demand is based on a load demand of at least one other functioning power system component in the power system.

According to another embodiment, a system for balancing load demands on power
20 system components comprises a first set of power system components in the power system and a load manager controlling load demands on the first set of the power system components based on a load balancing scheme.

According to yet another embodiment, an apparatus for managing load demands in a power system comprises means for determining whether load demands on one or more power

system components in the power system need to be varied; means for determining new load demands to be placed on the power system components in response to determining the load demands need to be varied; and means for controlling the load demands on the power system components to be substantially equal to the determined new load demands.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures in which like numeral references refer to like elements, and wherein:

Figure 1 illustrates a block diagram of a power system, according to an embodiment
10 of the invention;

Figures 2A-B illustrate tables providing examples of load balancing, according to embodiments of the invention;

Figure 3 illustrates a data flow diagram for the power system shown in figure 1, according to an embodiment of the invention;

15 Figure 4 illustrates a block diagram of a portion of a power system using a PUTS, according to an embodiment of the invention;

Figure 5 illustrates a flow chart of a method for managing load demands when a power system component fails, according to an embodiment of the invention;

20 Figure 6 illustrates a flow chart of a method for managing load demands when a power system is in a steady state, according to an embodiment of the invention;

Figure 7 illustrates a load management platform, according to an embodiment of the invention; and

Figure 8 illustrates conventional load sharing.

25

DETAILED DESCRIPTION OF EMBODIMENTS

According to an embodiment of the invention, a load manager is used to manage load demands on components of a power system. In one embodiment, when a failure of a power system component is detected, the load manager determines optimal load demands for the functioning power system components based on a load balancing scheme. In another embodiment, the load manager determines optimal load demands based on a load balancing scheme when the power system is in a steady state, such as when no power system component failures are detected. Load balancing may be performed to minimize the possibility of overloading a power system component; to minimize the possibility of the snowball effect; and to minimize costs due to over provisioning.

Figure 1 illustrates a power system 100, according to an embodiment of the invention. The power system 100, for example, may be used in a data center to supply power to a load, such as the computer systems 150a-l. The power system 100 is connected to a power utility grid 110 via a transfer switch 120. The power system 100 may also be connected to alternative energy sources, such as generators 112 and batteries 114. The transfer switch 120 controls which energy source is used to supply power to the power system 100. For example, the power utility grid 110 may be used as primary power source for the power system 100. If the power utility grid 110 fails or sufficient power is not being provided by the power utility grid 110, the transfer switch 120 supplies power to the power system 100 from the alternative energy sources. Alternatively, the alternative energy sources may be used as the primary power source for the power system 100, for example, because power may be supplied from the alternative energy sources at a cheaper rate. Then, power may be drawn from the power utility grid 110 as needed, for example, if the alternative energy sources cannot meet the load demand.

Other components of the power system 100 include UPSs 130a-d and PDUs 140a-f. The UPSs 130a-d are uninterruptible power sources that receive power from an energy source, such as the power utility grid 110, the generators 112 and/or the batteries 114. The UPSs 130a-d may provide uninterrupted power for at least a predetermined period of time to the load. For example, the UPSs 130a-d may supply uninterrupted power to the loads when the generators 112 are brought on line. Also, the UPSs 130a-d include circuits for minimizing undesired features of the power source, such as sags, surges, bad harmonics, etc.

The UPSs 130a-d are connected to the PDUs 140a-f. The PDUs 140a-f are power distribution units that supply power to the power supplies of the computer systems 150a-l, which may be housed in racks, such as the racks 160a-d. The PDUs 140a-f may include AC/AC power supplies, circuit breakers, power failure alarms, and other power conditioning circuits to step down the voltage and condition power supplied to the computer systems 150a-l. The computer systems 150a-l may include power supplies, not shown, that receive power from the PDUs 140a-f. The power supplies may be internal to the computer systems 150a-l or housed in the racks 160a-d.

Redundancy may be provided at one or more levels of the power system 100. The power system 100 provides N+1 redundancy, where N=1 at one or more levels. However, the power system 100 may also be provided with greater redundancy, e.g., 3 + 1, 2N+ 1, etc. The power system 100, also referred to as a grid, includes multiple levels in the grid. Each level may have N+1 redundancy. For 1+1 redundancy at the UPS level, each UPS 130a-d is connected to the transfer switch 120 using two separate electrical circuits (not shown) in the transfer switch 120 and two wires. Thus, the failure of any one circuit will not necessarily cause any of the computer systems 150a-l to lose power. Similarly, at the PDU level, each of the PDUs 140a-f is connected to at least two of the UPSs 130a-d. Thus, if for example the

UPS 130a fails, the UPS 130b supplies power to the PDU 140a. Redundancy may also be provided at the PDU level. For example, the computer system 150a may draw current via circuit 1 and circuit 2, where the circuits 1 and 2 are connected to two different power distribution circuits in the PDU 140a so there is no single point of failure. Also, the computer
5 system 150d receives current via circuits 3 and 4 connected to PDU 140c and PDU 140b respectively. Circuits 1-4 may include circuits in the PDUs or connected to the PDUs that distribute power to the loads. For example, the PDU 140a may include multiple power output channels, wherein each channel is connected to a load via a circuit breaker. A circuit may include a power output channel and/or other power circuits, such as a circuit breaker,
10 connected to the power output channel. At the computer system level, two power supplies may be used for each computer system to provide redundancy.

The power system 100 includes a load manager 160 for distributing the load demand placed on the components at each level of the power system 100. In a conventional power system, the load demand is shared approximately equally among redundant components to the
15 extent allowed by the discrete amount of power draw from the computer systems (i.e., the load). If one of the redundant power system components fails, the other power system component supports the entire load, which can lead to failure of other power system components and possibly complete power loss due to the snow ball effect. According to an embodiment, the load manager 160 distributes the load demand among substantially all the
20 power system components in a level of the power system according to a load balancing scheme to substantially minimize the possibility of overloading any one of the power system components in the level.

In one embodiment, the load manager 160 may implement load balancing at each level in the power system 100 by controlling the amount of current each component draws

from the next higher level. For example, the PDUs 140a-f are designed so that the load manager 160 directs each of the PDUs 140a-f to draw X% of its current from one of the UPSs 130a-d and (1-X)% from another one of the UPSs 130a-d, with the value of X controllable by the load manager 160. Similarly, the computer system 150a-l may each include at least two power supplies for redundancy, and the load manager 160 directs each power supply to draw X% of the load demand from one the PDUs 140a-f or circuits 1-4, and (1-X)% from another one of the PDUs 140a-f or circuits 1-4, with the value of X controllable by the load manager 160. For example, if circuit 1 fails and the load demand on circuit 2 increases as a result of the failure, the load manger 160 may direct the computer system 150d to reduce its loading of circuit 2 and increase its loading of circuit 3. Thus, in one embodiment, the load manager 160 may direct power system components (e.g., the computer systems 150a-l) at one level below the PDUs 140a-f to balance the load demand on the PDUs 140a-f. This embodiment may be performed at any level in the system 100, and is not limited to balancing the load demand at the PDU level. Load balancing is further illustrated by a table 200 shown in figure 2A.

The table 200 shown in figure 2A illustrates an example of load balancing implemented by the load manager 160, which can be used to reduce the amount of over provisioning and reduce the likelihood of the snow ball effect. The table 200 is described with respect to the power system 100 shown in figure 1. The circuits 1-4 supply power to the loads in the equipment racks 160a-b, such as the computer systems 150a-f, in an N+1 redundancy implementation, where N=1. As illustrated in row 1 of the table 200 and as illustrated by in figure 1, circuit 1 may support $\frac{1}{2}$ of the load demand of the computer system 150a, $\frac{1}{2}$ of the load demand of the computer system 150b, and $\frac{1}{2}$ of the load demand of the computer system 150c. Thus, the load demand on circuit 1 is shown as 1.5 ($\frac{1}{2} + \frac{1}{2} + \frac{1}{2}$) in row

1, column 203 of the table 200. The circuit loads in the column 203 represent equivalent loads on circuits 1-4, and may be used to compare load balancing according to an embodiment of the invention to conventional load sharing. An equivalent load unit may include the total load on circuits 1-4 divided equally among the circuits 1-4. As shown in row 5 2 of the table 200, circuit 2 supports $\frac{1}{2}$ of the load demand of each of the computer systems 150a, 150e, and 150f. Circuit 3 supports $\frac{1}{2}$ of the load demand of each of the computer systems 150d, 150e, and 150f, and circuit 4 supports $\frac{1}{2}$ of the load demand of each of the computer systems 150b, 150c, and 150d. Thus, in the steady state, when there are no component failures, the loading on each of the circuits 1-4 is 1.5 for the loads (e.g., computer 10 systems 150a-f).

In rows 1-4 of the column 204 of table 200, examples of loading on the circuits 1-4 are shown for two load sharing techniques. One technique is a conventional technique where the load demand is shared approximately equally by redundant components and upon failure of a component, 100% of the load is met by the functioning component. The second load 15 sharing technique is a load balancing scheme implemented by the load manager 160 according to an embodiment of the invention. With regard to the conventional load sharing technique, if circuit 1 fails, the load demand on the circuit 1 is zero. However, because the circuit 1 was responsible for 50% of the load demand of the computer systems 150a-c, the redundant circuit for the computer systems 150a-c now must support 100% of the load 20 demand of the computer systems 150a-c. This is shown in rows 2 and 4. The load demand on the circuit 2 due to the failure of circuit 1 is now 1 or 100% of the load demand of the computer system 150a, resulting in a total load demand on the circuit 2 of 2 equivalent load units. Likewise, the total load demand on the circuit 4 is now 2.5, because the circuit 4 must now meet 100% of the load demand of the computer systems 150b-c. Comparing the

conventional load sharing technique to the load balancing embodiment implemented by the load manager 160, the load manager 160 controls the load demand on the functioning circuits 2-4 to be substantially equal, such that the load demand of the computer systems 150a-f is shared by the circuits 2-4. For example, as shown in rows 2 and 4 of the table 200, the load demand on the circuit 3 includes 100% of the load demand of the computer system 150d, resulting in equal sharing of the load demand of the computer systems 150a-f. In other embodiments, the load manager 160 may divide the load demand based on other load balancing schemes, instead of providing equal sharing of the load demand. For example, if one circuit has a greater capacity, then that circuit may be directed to support a greater portion of the load demand.

Rows 5-8 of the table 200 show the load demands on the circuits 1-4 in column 203 if circuit 2 also fails after circuit 1 fails. As shown in rows 5 and 6, the load demands on circuits 1 and 2 are zero, because circuit 2 also failed. As a result, for the conventional load sharing technique the load demand on circuit 3 is 2.5 and the load demand on circuit 4 is 2.5. Assume that circuit 2 has a maximum load value rating that is approximately equal to 2.3 equivalent loads. In this scenario, the failure of circuits 1 and 2 using the conventional load sharing technique results in circuit 3 being overloaded, thereby bringing all but computer systems 150b-d off line. For the load balancing embodiment, the load manager 160 shifts the load demand of computer system 150d from circuit 3 to circuit 4, thereby maintaining the load demand on circuit 3 below 2.3. As shown in the example provided in the table 200, load balancing according to an embodiment of the invention reduces the maximum capacity that the components of the power system 100 must support. Consequently, the power system 100 may be provisioned with less over capacity, yielding a cost savings. Secondly, load balancing according to an embodiment of the invention may take into consideration the maximum

capacities (i.e., maximum load values) of power system components to minimize the possibility of overloading power system components.

Table 250 shown in figure 2B illustrates another example of load balancing whereby a UPS fails instead of a PDU circuit. Figure 2B is described with respect to the power system 100 shown in figure 1. The UPSs 130a-d supply power to the loads, e.g., PDUs 140a-f, to meet the load demand of the computer systems 150a-l. As illustrated in row 1, columns 252-253 of the table 250 and as illustrated by the connections of the UPSs 130a-d to the PDUs 140a-f shown in figure 1, the UPS 130a supports $\frac{1}{2}$ of the load demand of each of the PDUs 140a-c. Thus, the load demand on the UPS 130a is shown as 1.5 ($\frac{1}{2} + \frac{1}{2} + \frac{1}{2}$). The UPS 130b supports $\frac{1}{2}$ of the load demand of each of the PDUs 140a-b and 140d. The UPS 130c supports $\frac{1}{2}$ of the load demand of each of the PDUs 140c and 140e-f. The UPS 130d supports $\frac{1}{2}$ of the load demand of each of the PDUs 140d-f. Thus, in the steady state, when there are no component failures, the load demand on each of the UPSs 130a-d is 1.5.

In rows 1-4, columns 255-258 of the table 250, examples of load demands on the UPSs 130a-d are shown for two load sharing techniques. One technique is a conventional technique where the load demand is shared equally by redundant components and upon failure of a component, 100% of the load demand is met by the functioning component. The second load sharing technique is a load balancing scheme implemented by the load manager 160 according to an embodiment of the invention. With regard to the conventional load sharing technique, if the UPS 130a fails the load demand on the UPS 130a is zero. However, because the UPS 130a was responsible for 50% of the load demand of the PDUs 140a-c, the redundant UPSs for the PDUs 140a-c now must support 100% of the load demand of the PDUs 140a-c. This is shown in rows 2 and 3, column 255 of the table 250. The load demand on the UPS 130b after the failure of the UPS 130a is 2.5 resulting from the increased load

demand, i.e., 100% of the load demand, from each of the PDUs 140a and 140d. The load demand on the UPS 130c after the failure of the UPS 130a is 2 resulting from the increased load demand.

Comparing the conventional load sharing technique to the load balancing embodiment
5 implemented by the load manager 160, the load manager 160 controls components of the power system 100 such that the functioning UPSs 140b-d substantially equally share the load demand of the PDUs 140a-f. In one embodiment, the current drawn from a UPS by a PDU may be controlled by the PDU instead of the UPS. Thus, the load manager 160 may direct the PDUs 140a-f to vary the amount of current drawn from the functioning UPSs 130b-d for
10 load balancing.

As shown in rows 2-4, columns 257-258 of the table 250, the load manager 160 directs, for example, the PDUs 140a-b to draw current from the UPS 130b for supporting 100% of the load demand and directs the PDU 140d not to draw any current from the UPS 130b. Instead, the load manager 160 directs the PDU 140d to only draw current from the
15 UPS 130d, such as shown in row 4, column 258 of the table 250. This results in equal sharing of the load demand of the computer systems 150a-l by the functioning UPSs 130b-d, such as shown in column 257.

As described above, the load manager 160 may direct power system components at one level in the power system to increase or reduce current draw from power system
20 components one level higher to balance load demands on the power system components in the higher level. For example, the load manager 160 may direct the power systems of the computers 150a-f to increase or reduce power draw on the circuits 1-4 to control load demands on the circuits 1-4. Similarly, the load manager 160 may direct the PDUs 140a-d to increase or reduce power draw on the UPSs 130a-d to control load demands on the UPSs

130a-d. It will be apparent to one of ordinary skill in the art that other techniques may be used for balancing load demands. For example, the load manager 160 may direct the PDUs 140a-d to control their power output, such as the power output of the circuits 1-4, to balance the load demands on the PDUS 140a-d.

5 Figure 3 illustrates a data flow diagram for the power system 100 of figure 1, according to an embodiment of the invention. Figure 3 includes power system components 310 connected to the load manager 160. The power system components 310 may include the components of the power system 100 shown in figure 1, such as the transfer switch 120, the UPSs 130a-d, the PDUs 140a-f, the circuits 1-4, the computer systems (or the computer
10 system power supplies) 150a-l, etc. The power system components 310 may include sensors 312 for measuring the load demand on the power system components 310, i.e., the sensor data 314. The sensors 312 may include conventional power measurement circuits, such as current and/or voltage measuring circuits.

 The sensor data 314 is transmitted to the load manager 160 from the sensors 312.
15 The load manager 160 determines from the sensor data 314 whether the load demands on one or more of the power system components 310 need to be balanced. If the load demands need to be balanced, the load manager 160 sends load balancing control data 316 to the power system components 310. The load balancing control data 316 may include data associated with the amount of loading to be applied to power system components to balance the load
20 demand for those components. In one embodiment, the load balancing control data 316 is transmitted to the load (e.g., power system components at a lower level) instructing the load to reduce or increase current draw on the power system components (e.g., power system components at one or more levels above the load) to balance the load demand on those power system components. The load balancing control data 316 may include an amount of power to

be consumed by each of the loads. For example, the load manager 160 may direct the power supplies of one or more of the computer systems 150a-l to vary their load demand by drawing more or less current from the corresponding circuits connected to the PDUs 140a-f to balance the load demand on those circuits. Similarly, the load manager 160 may direct one or more of the PDUs 140a-f to vary their load demands on the UPSs 130a-d to balance the load demand on the UPSs 130a-d. In another embodiment, the load balancing control data 316 may be transmitted to the power system components having their load demands balanced, and may include the amount of power to be supplied to each load. For example, if the load demand of one of more of the PDUs 140a-f is to be varied, the load manager 160 may instruct, for example, the PDU 140b to decrease its power output of circuit 4 to a particular load and may instruct the PDU 140d to increase its power output of a circuit connected thereto for balancing the load demand on the PDUs 140a-f.

According to an embodiment, the load manager 160 may detect failure of one of the power system components 310 using the sensor data 314 and perform load balancing, such as shown in the examples provided in the tables 200 and 250 of the figures 2A-B. Failure detection, for example, may be determined by sensing no loading of one of the power system components 310 or by sensing overloading of one of the power system components 310.

The load manager 160 is connected to a load balancing repository 320 storing load balancing data. In one embodiment, the load manager 160 builds a model of the power system 100 and populates the model with a state of the power system 100 that exists if a failure is detected, assuming that load balancing is not performed. Then, the load manager 160 solves the model to determine the optimal load demands for the power system components 310 in view of the failure. The solution is saved in the load balance repository 320. This process is repeated, modeling different failures each time, and storing the solutions

(e.g., balanced load demands on the power system components 310) in the load balance repository 320. When an actual failure is detected by the load manager 160, based for example on the received sensor data 314, the load manager 160 queries the load balance repository 320, such as shown as load balance request 322, with the state of the power system 322, including the detected failure, to retrieve the load balancing solution, such as shown as load balance results 324, to be implemented in the power system 100.

According to other embodiments, when the power system 100 is in a steady state the load manager 160 may still invoke load balancing. The power system 100 may be in a steady state if no failures of power system components are detected. Also, even if a power system component has failed, but load balancing was performed and/or the failed component was replaced, the power system 100 may reach a steady state with no abnormal fluctuation of load demands. In one embodiment when the power system 100 is in a steady state, the load manager 160 may receive a load change request 330 to balance the load demand on one or more of the power system components 310. Such balancing may be used to allow maintenance on one of the power system components 310 or to free up capacity in a given set of power system components 310 to allow for the deployment of a new computer system. The load change request 330, for example, may be initiated by a system administrator (not shown).

In another embodiment when the power system 100 is in the steady state, the load manager 160 may monitor the load demands on the power system components 310 to determine whether load demands meet predetermined conditions. For example, the load manager 160 may be requested to maintain power system components 310 in a given level of the power system 100 to support substantially the same fraction of the total load demand of

the power system components 310 in the level. For example, each of the UPSs 130a-d source $\frac{1}{4}$ of the load demand on the UPSs 130a-d.

Another example of the predetermined conditions may include maintaining substantially the same spare capacity for power system components 310 in a given level. For example, the UPSs 130a-b may have twice the capacity of the UPSs 130c-d, and the load demands on the UPSs 130a-b may be adjusted to be twice the load demands of the UPSs 130c-d so as to provide equal tolerance for a failure. In yet another example, the predetermined conditions may be related to providing more tolerance for critical loads to decrease the possibility of power loss to the critical loads. In yet another example, the predetermined conditions may be related to maintaining the load demand on the power system components 310 below predetermined thresholds. The thresholds may be associated with a maximum capacity of the power system components, such as maximum power output, maximum load current, etc. The thresholds may include tolerances, such as being below a rated maximum capacity. Also, the thresholds may be related to a rate of increase of load demand. In addition to the load balancing solutions generated from modeling the power system 100 in different states, the load balance repository 320 may store the thresholds and other data needed to perform load balancing, which may be used in the steady state. For example, the load balance repository 320 stores data identifying the connection of a power system component 310 to other power system components 310. Thus, the load manager 160 may determine which loads need to be balanced based on which power system components 310 are able to service a load. Also, the load balance repository may store the sensor data 314 periodically received by the load manager 160.

The power system 100 may optionally include point-of-use transfer switches (PUTS) 340 for executing a quick transfer of load demand among the power system components 310.

Other fast load transfer devices performing the same function as the PUTS 340 may also be used. If load demands cannot be balanced by the load manager 160 within a predetermined period of time, such as due to software and data processing limitations, then PUTS 340 may be used to execute a fast transfer of loads (e.g., within 1/4 cycle or 4-6 ms) to predetermined sources to prevent overloading of one of the power system components 310. Then, the load manager 160 may rebalance the load demands to provide optimal loading of the power system components 310.

Figure 4 is a block diagram of a portion of the power system 100 including a PUTS 340 connected between the PDU 140 and the equipment rack 160a housing the computer systems 150a-c. In the steady state, circuits 1 and 2 supply power to the rack 160a. If circuit 2 fails, then the PUTS 340 switches the source from circuit 2 to circuit 4 connected to the PDU 140b and circuits 1 and 4 supply power to the rack 160a. Then, the load manager 160 balances the load demand on the functioning circuits connected to the PDUs 140a-f based on a load balancing scheme, which may include distributing some of the load demand on circuit 4 to other circuits. By using the PUTS 340, circuit 1 may not need to meet the load demand of the computer systems 150a-c housed in the rack 160a on its own. Thus, a lower capacity circuit may be used, reducing costs. In addition, the PUTS 340 allows fast switching of sources to substantially prevent overloading circuit 1. It will be apparent to one of ordinary skill in the art that multiple PUTS may be used in the power system 100, including multiple PUTS in figure 4. Furthermore, a PUTS may be incorporated in each rack 160a-d, instead of using multiple external PUTS connected to the PDU circuits, to minimize costs.

Figure 5 illustrates a method 500 for balancing load demand in a power system in response to detecting a failure of a power system component, according to an embodiment of the invention. The method 500 is described with respect to the power system 100 shown in

figures 1 and 2 by way of example and not limitation. Furthermore, the steps of the method 500 may be performed by software, hardware or a combination thereof.

At step 510, the load manager 160 detects a failure of a power system component of the power system 100. Failure detection may be based on the sensor data 314 received from the power system components 310 of the power system 100 shown in figure 2.

At step 520, the load manager 160 determines an optimal loading of components in the power system in view of the failed power system component. In one embodiment, the load manager 160 determines optimal load demands for the functioning components in the same level of the power system 100 as the failed component. For example, referring to figure 1, the PDU level includes the PDUs 140a-f. If circuit 1 of PDU 140a fails, then the load manager 160 determines the optimal load demand to be placed on the functioning PDUs 140b-f or circuits connected thereto. The power system components in one level may provide a level of redundancy for their respective loads. Thus, failure of one component in the level affects the load demand on at least one other component in the level, depending on the level of redundancy. As described above, in order to balance load demands at a target level of the system or grid, the load manager 160 may direct components at one level below the target level to increase or reduce current draw on the functioning components at the target level. Also, the load manager 160 may direct the functioning components at the target level to increase or reduce their power outputs to balance load demands at the target level instead of instructing the components at one level below the target level to increase or reduce current draw.

In one embodiment, the load manager 160 may determine optimal loading of the power system components based on previous modeling of failures in the power system 100. Different states of the power system 100 are modeled or simulated by the load manager 160.

Modeling of the different states includes determining resulting load demands on the power system components in response to one or more failed components, and determining optimal load demands that may be applied to the functioning components based on a load balancing scheme. Different states of the power system 100, including failures of different power
5 system components, are modeled to determine optimal load demands for the different states. The optimal load demands may then be used as a look-up table. Thus, when an actual failure is detected, optimal load demands are retrieved from the look-up table that are generated from a modeled state similar to a current state of the power system including the actual detected failure of a power system component. In other embodiments, the load manager 160 may
10 balance load demands by calculating load demands for the power system components based on a load balancing scheme being implemented and the measured current load demands.

The load manager 160 may use different load balancing schemes to determine optimal load demands for the power system components at step 520. Load balancing schemes may include substantially equally dividing a total load demand on functioning power system
15 components in a level where a failed power system component is detected, such as described with respect to Tables 200 and 250 in figures 2A-B. The load balancing scheme may take into consideration maximum loading values for the power system components. A maximum loading value, for example, may include a maximum rated power output or power demand that the power system component is designed to support. If the maximum loading value is
20 exceeded, the power system component may fail. Other load balancing schemes may include balancing load demands such that power system components have substantially equal spare capacity, or greater spare capacity may be provided for power system components servicing critical loads. It will be apparent to one of ordinary skill in the art that these and other load balancing schemes may be implemented by the load manager 160.

Fast load transfer devices, such as the PUTS 340 shown in figures 3 and 4, may optionally be used in the power system 100 to transfer load demand in minimal times to prevent overloading of a power system component before load balancing can be performed. If fast load transfer devices are activated due to overloading caused by the failed power system component, then, at steps 530 and 540, the load manager 160 controls the load demands on the power system components to be substantially equal to the load demands determined at step 520 based on the load balancing scheme. In one embodiment, the load manager 160 directs the power system components, such as the functioning power system components in the same level as the failed power system component, and the fast load transfer devices to balance the load demands on the power system components as determined at step 520. If fast load transfer devices are not used in the power system 100 or are not affected by the failed power system component, then the load manager 160, at steps 530 and 550, directs the power system components, such as the functioning power system components in the same level as the failed power system component, to balance the load demands on the components as determined at step 520. In one embodiment, the load manager 160 may control a load demand on a power system component by directing the power system to increase or reduce its power output to meet a load demand determined at step 520. In another embodiment, the load manager may direct a load on the power system component to increase or reduce its current draw on the power system component to vary the load demand on the power system component.

Figure 6 illustrates a method 600 for balancing load demand in a power system in a steady state. The method 600 is described with respect to the power system 100 shown in figures 1 and 2 by way of example and not limitation. Furthermore, the steps of the method 600 may be performed by software, hardware or a combination thereof.

At step 610, the load manager 160 determines whether load balancing is needed. Load balancing in the steady state may be performed for reasons other than detection of a failed power system component. For example, load balancing may be performed in response to a request received by the load manager 160. The load manager 160 may receive a request
5 330 to balance load demands on one or more of the power system components 310 shown in figure 3 to allow maintenance on one of the power system components 310 or to free up capacity in a given set of power system components 310 to allow for the deployment of a new computer system.

Also, the load manager 160 may balance load demands in the power system 100 to
10 meet predetermined conditions. For example, the load manager 160 may periodically monitor the load demands on the power system components 310. If the load demands fall out of balance, the load manager 160 balances the load demands, which may include varying one or more of the load demands based on a load balancing scheme being implemented by the load manager 160.

15 At step 620, if load balancing is needed, the load manager 160 determines load demands for power system components in the power system 100 based on the load balancing scheme being implemented at step 610. Different load balancing schemes may be implemented by the load manager 160. Load balancing schemes may include substantially equally dividing a total load demand on functioning power system components in a level
20 where a failed power system component is detected, such as described with respect to Tables 200 and 250 in figures 2A-B. The load balancing scheme may take into consideration maximum loading values for the power system components. A maximum loading value, for example, may include a maximum rated power output or power demand that the power system component is designed to support. If the maximum loading value is exceeded, the

power system component may fail. Other load balancing schemes may include balancing load demands such that power system components have substantially equal spare capacity, or greater spare capacity may be provided for power system components servicing critical loads.

It will be apparent to one of ordinary skill in the art that these and other load balancing
5 schemes may be implemented by the load manager 160.

Also, in one embodiment, the load manager 160 may model different states of the power system 100 to determine optimal load demands for different states of the power system 100 based on the load balancing scheme being implemented. The optimal load demands may be used to populate the data repository 320 shown in figure 3, which may include a look-up
10 table or database. The load manager 160 determines the current state of the power system from the measured load demands, for example, from the sensor data 314 shown in figure 3 received from the sensors 312. Using the current state, the load manager 160 queries the look-up table for optimal load demands to be applied to the power system components 310. In other embodiments, the optimal load demands may be periodically calculated from the
15 sensor data 314.

At step 630, the load manager 160 controls the load demands on the power system components 310 to apply the balanced load demands determined at step 620. In one embodiment, the load manager 160 directs the power system components to vary their load demands such that the load demands are substantially equal to the load demands determined
20 at step 620. In another embodiment, the load manager 160 may direct a load to increase or decrease its current draw from a source, i.e., a power system component that needs to change its load demand to a load demand determined at step 620, such that the load demand on the source is substantially equal to the load demand determined at step 620. For example, referring to figure 1, the load demand on the UPS 130a may be varied by directing the PDU

140a to vary the current draw on the UPS 130a, i.e., to vary the amount of power required by the UPS 130a from the PDU 140a.

One or more of the steps of the methods 500 and 600 may be repeated substantially continuously, periodically or demand driven to perform load balancing in the power system

5 100. These and other variations to the methods 500 and 600 will be apparent to one of ordinary skill in the art.

Figure 7 illustrates an exemplary platform 700 for the load manager 160. In one embodiment, the steps of the methods 500 and 600 performed by the load manager 160 may be performed by software or software in combination with hardware. The software may be
10 executed on the platform 700. The platform 700, for example, includes one or more processors, such as processor 702. Commands and data from the processor 702 are communicated over a communication bus 704. The platform 700 also includes a main memory 706, such as a random access memory (RAM), where the program code for the load manager 160 may be executed during runtime, and a secondary memory 708. The secondary
15 memory 708 includes, for example, one or more hard disk drives 710 and/or a removable storage drive 712, representing a floppy diskette drive, a magnetic tape drive, a compact disk drive, etc., where a copy of the program code for the load manager 160 may be stored. The removable storage drive 710 reads from and/or writes to a removable storage unit 714 in a well-known manner. User input and output devices may include a keyboard 716, a mouse
20 718, and a display 720. The display adaptor 722 interfaces with the communication bus 704 and the display 720 and receives display data from the processor 702 and converts the display data into display commands for the display 720. It will be apparent to one of ordinary skill in the arts that other known electrical components may be added or substituted in the platform 700. Also, one or more of the components in figure 7 may be optional (e.g., user input

devices, secondary memory, etc.). A network interface 724 may also be included for communicating with other computer systems and/or the power system components of the power system 100.

What has been described and illustrated herein are embodiments of the invention. The
5 terms, descriptions and figures used herein are set forth by way of illustration only and are not meant as limitations. Those skilled in the art will recognize that many variations are possible within the spirit and scope of the invention.